A Scalable Approach to Safety

Waabi's Voluntary Safety Self-Assessment

Version 1.1 (November 2023)

1. Introduction

Safety is paramount at Waabi. From day one, we've been committed to following industry best practices and meeting the highest standards for safe and responsible operation of autonomous vehicles.

Our voluntary safety self-assessment explains how we ensure safety throughout our development and in production. We believe that our products will only be perceived as safe if it is possible to understand why we are convinced that they are safe. That's why our voluntary safety self-assessment provides an overview of the elements we use to establish safe processes and how we make safe readiness decisions for supervised autonomy and driverless autonomy operations (see figure 1-1).

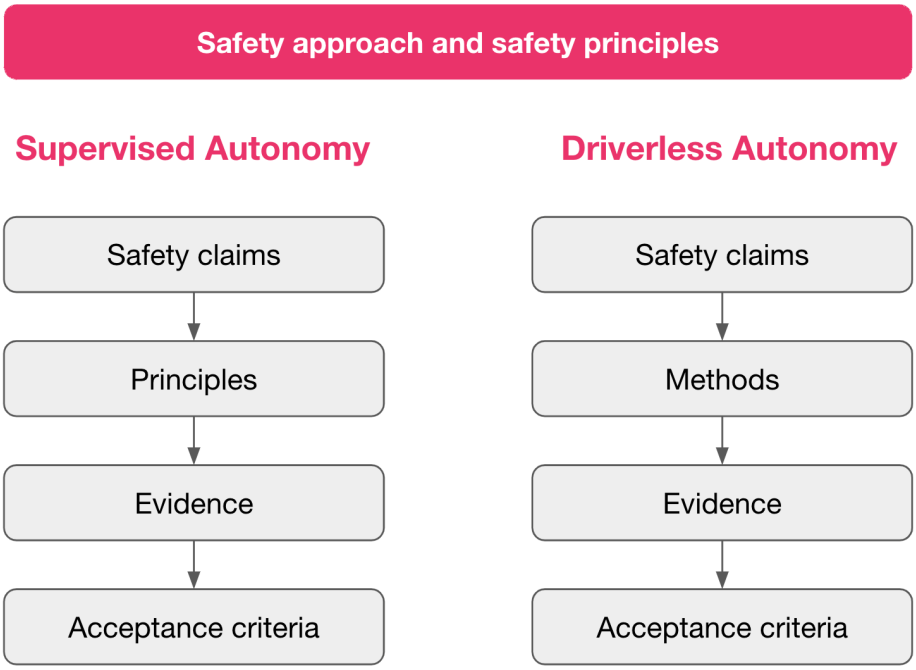

Figure 1-1: How Waabi's safety elements are leading to a readiness decision

This voluntary safety self-assessment will be a living document. Additional evidence and topics will be included in subsequent publications as well as our safety case.

Andre Strobel, Director of Safety Engineering at Waabi

2. Waabi Driver and Waabi World

Waabi World is the next generation high-fidelity, closed-loop end-to-end simulator. It is a realistic, controllable, diverse, fast, immersive and reactive simulation environment that automatically provides feedback to the Waabi Driver. Scenarios can be created by the user for specific situations, as digital twins of real world logs, using generative models to ensure broad and deep coverage as well as through adversarial optimization to identify, cover and solve the really hard problems.

Waabi is the only company building an AI-first solution to autonomous driving. The Waabi Driver, the next generation of autonomous trucking technology, is built on years of innovation in AI and learns from data on its own, enabling it to execute the complex decision-making needed for operating on the road safely.

Rather than relying solely on ineffective and inefficient real-world testing, we leverage Waabi World's realistic simulation to more quickly and more reliably perform a greater amount of testing in a safer and more efficient manner. By the time one of our autonomous trucks hits the road, it has already experienced countless driving situations in Waabi World, enabling it to safely navigate any scenario it might encounter, from typical driving conditions to increasingly challenging and unique incidents. On-road testing at Waabi is primarily reserved for the final validation (see figure 2-1).

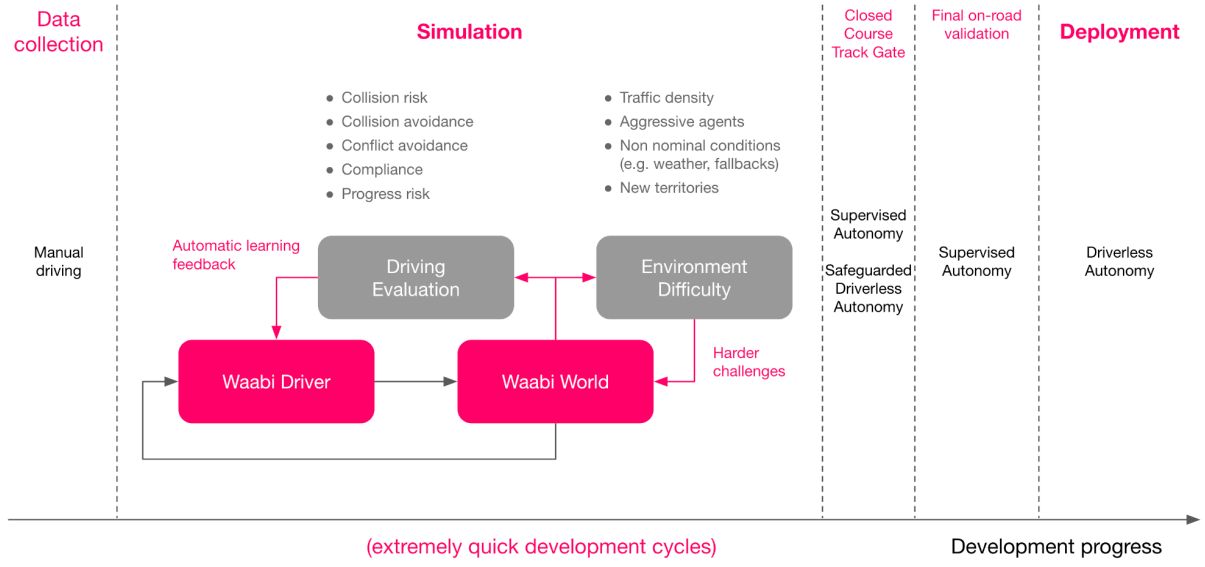

Figure 2-1: Waabi Driver in the Waabi World

3. Waabi's safety organization

3.1 Safety principles

Our unique approach to safety starts with Waabi World - our realistic simulation environment used to train and test the Waabi Driver in driving as safely as possible while fulfilling the driving mission.

Safety is not relative	*Our ability to develop safe driving skills is not limited by a human reference. We can exceed the performance of human drivers by learning while driving more miles than any human can in the real world during their lifetime.*
Safety is dynamic	*We deploy the Waabi Driver in a responsible way and continue to evolve its capabilities while being operational. Deploying responsibly includes all aspects of autonomous driving - our nominal driving skills in the real world, our reaction to faults as well as our operational policies and procedures that ensure we are welcome everywhere.*
Safety must be credible	*We understand that the performance of the Waabi Driver in the real world highly depends on how accurately Waabi World represents reality. Therefore, we take great care to measure and minimize any differences. Where we cannot cover hardware related aspects in simulation we conduct thorough bench and real world testing to ensure all features work as intended.*

Based on our approach to safety, the principles in table 3.1-1 guide us during our development and day-to-day operations.

Safe product	Safe workplace	Safe deployment
We see safety as the most important and ever evolving aspect of our driving.	We are proud of our safety practices in the workplace and our safety processes during development and production that are based on industry best practices and experience.	We release and deploy our driverless product in a credible and transparent way so others can follow our decision making.
We strive to be at least as safe as a licensed and unimpaired human driver in the same operational design domain (ODD) using the same base vehicle.		We follow known safety standards and help evolve the industry where existing standards are not sufficient.
We strive to improve safety for each road user that we interact with.	We listen to all concerns of our employees and the public in regards to the safety of our product.	Our vehicles can reach a minimal risk condition in the event of any single independent fault or any single set of dependent faults.
We will comply with the rules of the road more strictly than human drivers typically do.		Our vehicles will not disturb the normal flow of traffic in a significant way.

Table 3.1-1: Waabi safety principles

3.2 Governance of safety topics and escalation

A truly safety oriented organization has a clearly defined governance structure for how safety issues are escalated and resolved, who makes safety decisions, who prepares safety decisions and who provides the necessary information. At Waabi, the Waabi Safety Committee makes all safety decisions. The following are the responsibilities of the Waabi Safety Committee:

- Decide on the governance of product safety topics
- Decide on the prioritization of the safety case development
- Delegate development sign-offs to individuals in the organization
- Ultimately sign off the product

Every employee at Waabi can escalate a product or process related safety concern either directly through their manager, by directly contacting any person working in the safety area or through the Waabi safety concern form. The submission of the Waabi safety concern form is anonymous unless the employee adds their contact information.

3.3 Supervised autonomy and driverless autonomy

In the following, a highly autonomous vehicle is defined as a vehicle that has all the hardware and software that is required to drive by itself without any remote autonomy connection during its operation like assistance from humans.

If the highly autonomous vehicle is an upfitted base vehicle from an OEM that is typically driven manually by humans, the base vehicle is converted into an autonomy enabled vehicle with actuators and interfaces that allow adding the autonomy system like sensors, compute and software to finish the highly autonomous vehicle.

The highly autonomous vehicle can be operated in different modes that describe the status of the autonomy system and who is ultimately in control (see table 3.3-1).

Operation mode	Autonomy system status	Who is in control?	Purpose
Transport	Deactivated	Human driver	Manual transport outside of approved ODD and scope
Data Collection	Passive	Vehicle operator supported by a fleet specialist	Manual data collection
Supervised Autonomy	Active	Vehicle operator supported by a fleet specialist	Testing the highly autonomous vehicle (advanced SAE Level 2)
Driverless Autonomy	Active	Autonomy system	Normal highly autonomous vehicle operation (SAE Level 4)

Table 3.3-1: Terminology for the operational modes of an autonomous vehicle

The comparison to SAE Levels is done based on “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104” [9].

Each mode requires a different safety case. *Transport* and *Data Collection* modes are the manual driving modes for which the only requirement is that any present autonomy system cannot accidentally take over control of any driving task of the highly autonomous vehicle.

Supervised Autonomy mode is used during the development of the highly autonomous vehicle. In this mode the vehicle operator and fleet specialist are always able to safely take over control of the highly autonomous vehicle. The vehicle operator is sometimes also called a “safety driver”. The fleet specialist is inside the vehicle and monitors the autonomy system as well as the environment. The vehicle operator and fleet specialist continuously exchange information about the autonomy system and the environment. We call both of them autonomous specialists.

Driverless Autonomy mode is how the product will operate when it is deployed without a human on board who could control the driving tasks. In this mode the autonomy system is always in control and will make the ultimate driving decisions.

3.4 Safety processes

We follow industry standards where they apply and are useful to ensure the credibility of our approaches. For this we borrow suitable elements from the following standards:

- We follow the safety claim approach as described in UL 4600 [11][12]
- Our artifacts for supervised autonomy cover the content in SAE J3018 for “Guidelines for Safe On-Road Testing of SAE Level 3, 4, and 5 Prototype Automated Driving Systems (ADS)” [8]
- We apply hazard and risk analysis (HARA) as described in ISO 26262 [3] and follow this standard where it ensures a safe and reliable platform for our highly autonomous vehicle
- We use Systems Theoretic Process Analysis (STPA) [6] and consider ISO/PAS 21448 “Road vehicles - Safety of the intended functionality” [4] to avoid hazards
- We rely on the standards that the base vehicle OEM used to certify their product
- We also use general automotive, aerospace, industrial and medical standards when we design our hardware

We apply a thorough development process that starts with requirements that are implemented and then rigorously verified and validated.

As part of the verification and validation we assess the impact of all failed requirements to identify what changes are needed before we can sign off the desired scope.

For our safety approach we considered the 12 safety elements that are proposed by NHTSA in “Automated Driving Systems 2.0: A Vision for Safety” [7]:

1. System Safety
2. Operational Design Domain
3. Object and Event Detection and Response
4. Fallback (Minimal Risk Condition)
5. Validation Methods
6. Human Machine Interface
7. Vehicle Cybersecurity
8. Crashworthiness
9. Post-Crash ADS Behavior
10. Data Recording
11. Consumer Education and Training
12. Federal, State, and Local Laws

All of these safety elements are addressed as needed. The methods described in section 6.2 Driverless autonomy readiness methods focus on system safety, operational design domain, object and event detection and response, fallback (minimal risk condition), validation methods, human machine interface, vehicle cybersecurity, post-crash ADS behavior as well as federal, state and local laws. In several aspects they go beyond these safety elements and provide a larger framework that reaches from testing all the way to product deployment and sustained operation.

We are also considering the 13 safety assessment expected outcomes that are proposed by Transport Canada (🇨🇦) in “SAFETY ASSESSMENT for Automated Driving Systems in Canada” [10] which overlap with the NHTSA (🇺🇸) safety elements:

1. ADS Level of Automation and Intended Use 🇨🇦
2. Operational Design Domain 🇨🇦🇺🇸
3. Object and Event Detection and Response 🇨🇦🇺🇸
4. International Standards and Best Practices 🇨🇦
5. Testing and Validation 🇨🇦🇺🇸
6. Safety Systems 🇨🇦🇺🇸
7. Human-Machine Interface and Accessibility of Controls 🇨🇦🇺🇸
8. Public Education and Awareness 🇨🇦🇺🇸
9. User Protections during Collisions or System Failures 🇨🇦🇺🇸
10. Cyber Security 🇨🇦🇺🇸
11. System Updates and After-Market Repairs/Modifications 🇨🇦
12. User Privacy 🇨🇦
13. Collaboration with Government Agencies and Law Enforcement 🇨🇦

All of these safety assessment guidelines are addressed as needed. Our driverless autonomy product is an SAE level 4 system. During development we also use supervised autonomy which adds autonomous specialists as human supervisors to our product and hence degrades these prototypes to something closer to an SAE level 2 system. As mentioned in this section we follow international standards and best practices as they apply. In addition to the before mentioned safety elements the methods described in section 6.2 Driverless autonomy readiness methods also focus on system updates and after-market repairs/modifications, user privacy as well as collaboration with government agencies and law enforcement.

4. Waabi's safety approach

4.1 Absence of unreasonable risk

When driving in the real world there will always be some risk of causing harm and being harmed. A Waabi vehicle's behavior cannot eliminate the risk of other road users being reckless and initiating conflict. Also, a critical component of the base vehicle can malfunction unexpectedly, e.g. a steer tire blowout can make it hard for a human driver and a highly autonomous vehicle to stay safe. Therefore, we acknowledge that we will have to account for and mitigate risks that are outside of our control. In order to establish acceptance criteria for the readiness of our deployments, we prove the absence of unreasonable risk through our credible safety claims [1].

Risk can be defined as a function of severity and exposure [3]. We ensure the scope of our deployments so that the severity and the exposure are limited to situations for which we have the evidence that the *critical events*, *likely events* and *reasonably expectable events* are either *controllable events* or *tolerable events*.

Critical events are all safety relevant conflicts and all types of collisions. For example, a close cut-in in front of our highly autonomous vehicle is a *critical event*. *Likely events* occur in the selected ODD at a rate with which they are likely to occur at the desired exposure. For example, when driving 10,000 miles it is very likely to have more than one safety relevant conflict like a cut-in. *Reasonably expectable events* need to be considered as part of the safety case based on reasonable professional knowledge about the highly autonomous vehicle systems and ODD. For example, unknown hazards are not *reasonably expectable events*. Also, extremely *unlikely events* are also not *reasonably expectable events*.

By ensuring all *critical events*, *likely events* and *reasonably expectable events* are *controllable events* or *tolerable events* we also minimize the unreasonable risk. *Controllable events* are situations where the highly autonomous vehicle can use its actuators to avoid a conflict or collision. For example, a slower vehicle in front of the highly autonomous vehicle becomes a *controllable event* by applying the brakes and reducing speed accordingly. *Tolerable events* have either mitigations or won't affect the overall safety assessment for the desired exposure. For example, certain weather conditions can be avoided by reviewing weather reports upfront or by slowing down or stopping in case the weather changes unexpectedly.

The evidence that we provide as part of our deployment sign-offs for our safety claims in sections 5.1 Safety claims for supervised autonomy and 6.1 Safety claims for driverless autonomy cover all relevant aspects below the maximum severity and maximum exposure limits. We can set these limits by setting the scope of our deployments. For example, we can limit the goods we transport to not being hazardous materials and we can limit the exposure of a certain deployment to a maximum number of miles. In later deployments we can add these features to our scope as we scale. The relevant aspects consist of *frequent events*, *challenging events*, *rare events* and *high risk events* (see figure 4.1-1).

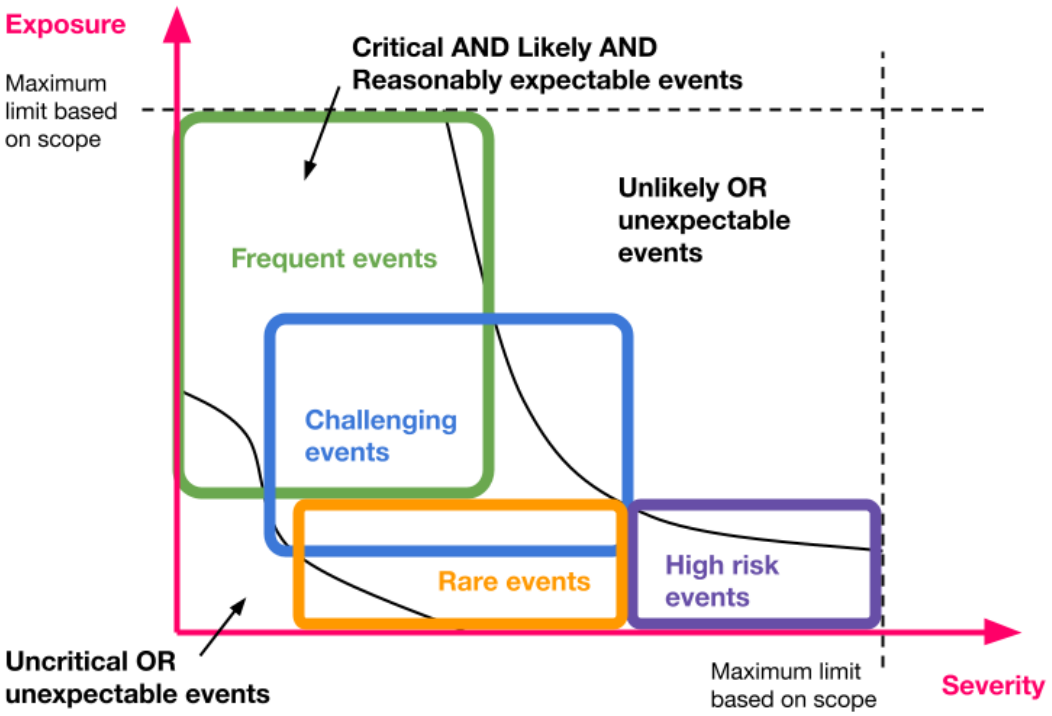

Figure 4.1-1: Focus areas when covering the risk space

Our vehicles, roadways and road rules are designed in a way that ensures events with a high severity don't happen at a high frequency, i.e. high severity and high exposure events are actually *unlikely events*. By covering the *high risk events*, we will also cover them if they start to appear more frequently.

4.2 Credibility of the assessment

A safety case for an autonomous vehicle service is only valid if the related safety claims are credible and proven [2]. The following guidelines help us to identify whether our set of safety claims is creating a credible and proven safety case for an autonomous vehicle service:

1. The safety claims must create a credible argument
    a. The safety claims cover all aspects of the autonomous vehicle service
    b. The safety claims can be measured and credible acceptance criteria exist
2. The evidence for the safety claims must show the absence of unreasonable risk
    a. The evidence covers all aspects of the safety claims
    b. The evidence proves that the overall risk is manageable
3. The quality of the safety claims and evidence is ensured
    a. The quality is regularly assessed by experts
    b. The quality is continuously improved by e.g. new knowledge or experiences

This document is not intended to prove the credibility for a particular sign-off. Instead, it is intended to show which aspects we consider while building our safety case [5].

4.3 Autonomous vehicle technology

The Waabi Driver perceives its environment using a variety of sensors that feed our autonomy system. It localizes itself in space and relative to other objects and then plans the best path forward considering the movement of the road users next to, ahead of and behind it. The desired path is then executed using precise motion control that commands steering, braking and acceleration (see figure 4.3-1).

Figure 4.3-1: Highly autonomous Waabi truck

Waabi is developing the Waabi Driver initially for heavy duty trucks that primarily operate on interstate highways. Driving at highway speeds with an up to 80,000 lbs heavy vehicle that is articulated and about 70 ft long requires very forward looking and defensive driving behaviors.

We use and create highly reliable components to minimize the likelihood of failures. We safeguard our system with a layered approach to prevent failures from becoming hazards through defense in depth, e.g. we check our actuation commands in several steps for reasonable values within our approved limits before sending them to the actuators.

We also add systems to our highly autonomous vehicle that simplify the monitoring and safe operation in the field.

5. Supervised autonomy sign-off

5.1 Safety claims for supervised autonomy

For supervised autonomy missions, the safety claims are based on having autonomous specialists who can always control our highly autonomous vehicles:

> *1* *We operate our autonomous vehicles with autonomous specialists who are always in control of the vehicle.*

We reason that this safety claim is fulfilled by the following subclaims:

> *1.1* *The base vehicle functions have been certified for the safe operation of the vehicle.*

> *1.2* *Our autonomy system and our operational procedures allow the autonomous specialists to always observe the behavior of the autonomy system and safely take over control with the base vehicle functions if needed.*

These subclaims have been decomposed into a comprehensive set of architectural, behavioral and operational principles (see section 5.2 Supervised autonomy principles) [13].

5.2 Supervised autonomy principles

The following architectural principles lay the foundation for a safe operation of the vehicle during supervised autonomy missions (see table 5.2-1).

Number	Architectural principle
A1	The base vehicle functionality and base vehicle performance have not been altered by the integration of the autonomy system in a way that it affects the safe operation of the vehicle. Hence, the autonomous specialists can use the base vehicle interfaces to take over control of the vehicle.
A2	The autonomous specialists can safely engage the supervised autonomy mode. Also, the autonomous specialists can always safely take over control by disengaging the supervised autonomy mode. When not in supervised autonomy mode the vehicle can be operated in the manual data collection mode.
A3	The autonomy system always informs the autonomous specialists whether the autonomy system is active or not and whether there is a fault and the autonomous specialists need to take over control.
A4	The vehicle operator will be informed about upcoming maneuvers of the highly autonomous vehicle.
A5	The vehicle operator can always control the vehicle in a conventional manner after the supervised autonomy mode has been disengaged.
A6	The autonomy system is designed in a way that the platform sign-off is valid for all vehicles of this type.

Table 5.2-1: Architectural principles

Behavioral principles are an additional layer of safety that makes it easier for the autonomous specialists to assess certain situations (see table 5.2-2).

Operational principles ensure everything is executed as intended and what to do if something doesn’t play out as intended (see table 5.2-3).

Number	Behavioral principle
B1	The autonomy system has sufficient basic behavioral competencies to ensure the autonomy system can execute simple maneuvers without the need for immediate or frequent disengages.
B2	The behavioral capability limits of the autonomy system have been assessed. A list of policy disengages has been created which describes when the system should not be operated in the supervised autonomy mode as it might be outside of its proven and safe behavioral capability limits.

Table 5.2-2: Behavioral principles

Number	Operational principle
O1	The autonomous specialists know how to safely operate the vehicle. This includes all procedures including the pre-trip inspection, special preparations for supervised autonomy missions, defensive driving behaviors, what to do in case of an incident during the mission as well as the post-trip inspection and report.
O2	The autonomous specialists can always decide to disengage. They are also trained on the list of policy disengages.
O3	The autonomous specialists regularly assess whether they are fatigued and do not continue to operate the vehicle if they feel any evidence of upcoming fatigueness.
O4	A dispatch role is assigned to a person in the organization. The dispatcher has the necessary information about what vehicles can be assigned to which missions and what restrictions currently apply. The dispatcher assigns the mission and vehicle to the autonomous specialists.
O5	Safety relevant events are being investigated and a field safety process is established.
O6	Procedures for who has to do what in case of an incident exist.
O7	All regulatory requirements for operating in manual data collection and supervised autonomy modes are met during normal operation and in case of incidents.
O8	All relevant regulatory authorities are informed upfront about our intentions to conduct supervised autonomy missions.

Table 5.2-3: Operational principles

5.3 Supervised autonomy sign-off evidence

As part of the supervised autonomy sign-off, we conduct a thorough architectural review and hazard analysis to identify the architectural requirements. This includes a detailed review and test of all the requirements and their implementation related to mode transitions within the compute, the actuators, and the related human-machine-interface concept. We also consider the controllability of the vehicle during these mode transitions. All architectural requirements are then rigorously tested.

Behavioral analysis is done based on simulation and closed course testing. Any deficiencies are then excluded from the scope by policy takeovers.

Table 5.3-1 shows what evidence is collected to cover the operational principles.

Description	O1	O2	O3	O4	O5	O6	O7	O8
Vehicle operator handbook	X		X		X	X		
Pre-trip inspection procedure and protocol	X							
Expected driving behaviors policy	X							
Disengage policy	X	X						
Incident response policy	X				X	X	X	
Vehicle operator monitoring procedure			X					
Vehicle operator and fleet specialist trainings	X	X	X			X		
Dispatcher role definition and assignment			X	X	X			
Tracking of the approved hardware, firmware and software combinations by missions type				X				
Summary of licenses and regulatory approvals							X	
Summary of upfront announcements to regulatory authorities								X
Field safety process				X	X	X		
Waabi safety concern form	X		X		X			

Table 5.3-1: Artifacts fulfilling the operational principles

In addition, we conduct operational readiness testing with which we assess whether all our operational procedures are understood and executed as intended. Any deficiencies lead to continuous improvement of our operational procedures.

6. Driverless autonomy sign-off

6.1 Safety claims for driverless autonomy

We see safety as the most important and ever evolving aspect of our driving, especially during driverless autonomy missions where the highly autonomous vehicle is ultimately in control at all times. We believe that technology can continuously improve traffic safety, which is captured in our first safety claim:

> 1 *Our autonomous vehicle operation contributes to improving traffic safety for everyone as we grow our service.*

We see making progress along our trip, and not stopping for every immaterial fault, as an important aspect of safety. We develop a system that can handle every to be expected fault without an immediate impact on progress which is captured in our second safety claim:

> 2 *Our autonomous vehicle service is as reliable as or more reliable than a fleet of human driven vehicles in all safety relevant aspects of our ODD.*

Experts are still an important part of our autonomous vehicle service. We operate our autonomous vehicle service responsibly during both development and when in production, which is captured in our third safety claim:

> 3 *Our autonomous vehicle operation is guided by experts who follow safe, robust and well established processes, policies and procedures.*

We want to find and learn from the hard and long tail issues quicker than having to drive a lot of miles in the real world. We use a highly realistic simulation-based approach to qualify our autonomy system which is captured in our fourth safety claim:

> 4 *Our autonomous vehicle is thoroughly tested and evaluated in full system simulation and the real world which is highly representative of our ODD, adversarial edge cases and the collective experience of some of the most experienced human drivers.*

We are developing more detailed subclaims for each of the driverless safety claims. These will be covered with methods that apply full system simulation and real world testing to provide the evidence that all driverless safety claims are fulfilled.

6.2 Driverless autonomy readiness methods

The evidence for our driverless autonomy sign-off is created by a variety of different methods that cover all relevant aspects of our autonomous vehicle service and the associated risks.

We define the targets for the acceptance criteria of these methods based on the exposure that we intend to sign off.

The verification and validation procedures and methods (shown in the boxes on the left) also address the safety elements (🇺🇸🇨🇦) and additional safety assessment expected outcomes (🇨🇦) that are described in section 3.4 Safety processes as well as further aspects (shown on the right):

Governance, field safety process and risk management	• *Methods that enable a holistic safety decision* • *System Updates and After-Market Repairs/Modifications* 🇨🇦
Conflict avoidance, safety relevant conflicts, collision avoidance and collision risk	• *System Safety* 🇺🇸🇨🇦 • *Object and Event Detection and Response* 🇺🇸🇨🇦
Compliance risk	• *Federal, State/Provincial and Local Laws* 🇺🇸🇨🇦 • *User Privacy* 🇨🇦
Progress risk	• *Method that addresses the negative impact of stops*
Non nominal conditions, incident management and fault management	• *Operational Design Domain* 🇺🇸🇨🇦 • *Post-Crash ADS Behavior* 🇺🇸🇨🇦 • *Fallback (Minimal Risk Condition)* 🇺🇸🇨🇦 • *Collaboration with Government Agencies and Law Enforcement* 🇨🇦
Comprehensive verification	• *System Safety* 🇺🇸🇨🇦 • *Object and Event Detection and Response* 🇺🇸🇨🇦 • *Human Machine Interface* 🇺🇸🇨🇦 • *Vehicle Cybersecurity* 🇺🇸🇨🇦

As we progress towards our first driverless autonomy sign-off, we will publish details about the key methods that will drive our readiness decision, along with the supporting research that is done at Waabi. These publications will focus on the performance of our simulation tools (step 1), how we evaluate the behaviors of the Waabi Driver in a scalable way (step 2), and the comprehensive verification and validation of our highly autonomous vehicle platform (step 3).

Most recently we published papers that lay the foundation for step 1 (see table 6.2-1).

Paper	Abstract
UniSim: A Neural Closed-Loop Sensor Simulator	Rigorously testing autonomy systems is essential for making safe self-driving vehicles (SDV) a reality. It requires one to generate safety critical scenarios beyond what can be collected safely in the world, as many scenarios happen rarely on our roads. To accurately evaluate performance, we need to test the SDV on these scenarios in closed-loop, where the SDV and other actors interact with each other at each timestep. Previously recorded driving logs provide a rich resource to build these new scenarios from, but for closed-loop evaluation, we need to modify the sensor data based on the new scene configuration and the SDV's decisions, as actors might be added or removed and the trajectories of existing actors and the SDV will differ from the original log. In this paper, we present UniSim, a neural sensor simulator that takes a single recorded log captured by a sensor-equipped vehicle and converts it into a realistic closed-loop multi-sensor simulation. UniSim builds neural feature grids to reconstruct both the static background and dynamic actors in the scene, and composites them together to simulate LiDAR and camera data at new viewpoints, with actors added or removed and at new placements. To better handle extrapolated views, we incorporate learnable priors for dynamic objects, and leverage a convolutional network to complete unseen regions. Our experiments show UniSim can simulate realistic sensor data with small domain gap on downstream tasks. With UniSim, we demonstrate, for the first time, closed-loop evaluation of an autonomy system on safety-critical scenarios as if it were in the real world.
MixSim: A Hierarchical Framework for Mixed Reality Traffic Simulation	The prevailing way to test a self-driving vehicle (SDV) in simulation involves non-reactive open-loop replay of real world scenarios. However, in order to safely deploy SDVs to the real world, we need to evaluate them in closed-loop. Towards this goal, we propose to leverage the wealth of interesting scenarios captured in the real world and make them reactive and controllable to enable closed-loop SDV evaluation in what-if situations. In particular, we present MixSim, a hierarchical framework for mixed reality traffic simulation. MixSim explicitly models agent goals as routes along the road network and learns a reactive route-conditional policy. By inferring each agent's route from the original scenario, MixSim can reactively re-simulate the scenario and enable testing different autonomy systems under the same conditions. Furthermore, by varying each agent's route, we can expand the scope of testing to what-if situations with realistic variations in agent behaviors or even safety critical interactions. Our experiments show that MixSim can serve as a realistic, reactive, and controllable digital twin of real world scenarios.

Table 6.2-1: Research papers related to step 1 of our driverless autonomy publications

A. Reference information

A.1 Abbreviations

Table A.1-1 lists all of the abbreviations used in this document.

Abbreviation	Meaning	Description
ADS	Autonomous Driving System	The system that drives a vehicle autonomously
AI	Artificial Intelligence	The methods of a computer system that performs a task that normally requires human intelligence
HARA	Hazard And Risk Analysis	The method to identify malfunctions that could lead to hazards, to rate the relevant risks of hazards and to formulate safety goals
ISO	International Organization for Standardization	An independent, non-governmental international organization with a membership of 168 national standards bodies
LiDAR	Light Detection And Ranging	A remote sensing method that uses light in the form of a pulsed laser to measure ranges
NHTSA	National Highway Traffic Safety Administration	A department of the United States Department of Transportation with the mission to save lives, prevent injuries, and reduce economic costs due to road traffic crashes, through education, research, safety standards and enforcement
ODD	Operational Design Domain	The sign-off scope in which the highly autonomous vehicle will be operating
OEM	Original Equipment Manufacturer	The manufacturer of the base vehicle and autonomy enabled vehicle
PAS	Publicly Available Specification	A standardization document that closely resembles a formal standard in structure and format, but has the objective to speed up standardization often in response to an urgent market need
SAE	SAE International	A United States-based, globally active professional association and standards developing organization for engineering

Abbreviation	Meaning	Description
		professionals in various industries
SDV	Self-Driving Vehicle	A vehicle that doesn't require a human driver
STPA	Systems Theoretic Process Analysis	A systems approach to hazard analysis based on the premise that accidents happen when we lose control, i.e. considering the analysis as a control problem and not as a failure problem

Table A.1-1: Abbreviations used and defined in this document

A.2 Literature

[1] Favarò, F. : *Exploring the Relationship Between "Positive Risk Balance" and "Absence of Unreasonable Risk"*. Retrieved from: https://arxiv.org/ftp/arxiv/papers/2110/2110.10566.pdf (2021)

[2] Favarò, F.; Fraade-Blanar, L.; Schnelle, S.; Victor, T.; Peña, M.; Engstrom, J.; Scanlon, J.; Kusano, K.; Smith, D. : *Building a Credible Case for Safety: Waymo's Approach for the Determination of Absence of Unreasonable Risk*. Retrieved from: https://arxiv.org/abs/2306.01917, arXiv:2306.01917 (2023)

[3] International Organization for Standardization : *Road vehicles - Functional safety - Part 1: Vocabulary. (ISO 26262-1:2018)*. Retrieved from: https://www.iso.org/standard/68383.html (2018)

[4] International Organization for Standardization : *Road vehicles - Safety of the intended functionality. (ISO/PAS 21448:2022)*. Retrieved from: https://www.iso.org/standard/77490.html (2022)

[5] Koopman, P. : *UL 4600: What to Include in an Autonomous Vehicle Safety Case*. IEEE Computer. Retrieved from: https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10109263 (2023)

[6] Leveson, N. G.; Thomas, J. P. : *STPA handbook*. Retrieved from: https://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf (2018)

[7] National Highway Traffic Safety Administration : *Automated Driving Systems 2.0: A Vision for Safety*. Retrieved from: https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/13069a-ads2.0_090617_v9a_tag.pdf (2017)

[8] SAE International : *Guidelines for Safe On-Road Testing of SAE Level 3, 4, and 5 Prototype Automated Driving Systems (ADS) J3018_201503*. Retrieved from: https://www.sae.org/standards/content/j3018_201503/ (2015)

[9] SAE International : *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104*. Retrieved from: https://www.sae.org/standards/content/j3016_202104/ (2021)

[10] Transport Canada : SAFETY ASSESSMENT for Automated Driving Systems in Canada. Retrieved from: https://publications.gc.ca/collections/collection_2019/tc/T86-52-2018-eng.pdf (2019)

[11] Uber : *A principled approach to safety*. Retrieved from: https://uber.app.box.com/v/UberATGSafetyReport (2020)

[12] UL Standards & Engagements : *ANSI/UL 4600 Standard for Safety for the Evaluation of Autonomous Products*. Retrieved from: https://ulse.org/ul-standards-engagement/presenting-standard-safety-evaluation-autonomous-vehicles-and-other-1 (2022)

[13] Webb, N.; Smith, D.; Ludwick, C.; Victor, T.; Hommes, Q.; Favarò, F.; Ivanov, G.; Daniel, T. : *Waymo's safety methodologies and safety readiness determinations*. Retrieved from: https://arxiv.org/abs/2011.00054, arXiv:2011.00054 (2020)